



US 20050149583A1

(19) **United States**

(12) **Patent Application Publication**
Baskaran et al.

(10) **Pub. No.: US 2005/0149583 A1**

(43) **Pub. Date: Jul. 7, 2005**

(54) **CUSTOMIZABLE METADATA MERGING FRAMEWORK**

Publication Classification

(76) Inventors: **Rajagopal Baskaran**, Palo Alto, CA (US); **Jaime E. Singson**, San Francisco, CA (US)

(51) **Int. Cl.⁷ G06F 12/00**

(52) **U.S. Cl. 707/203**

(57) **ABSTRACT**


One embodiment of the present invention provides a system that facilitates merging different versions of a database object. During operation, the system receives metadata associated with a first version of a database object and a second version of the database object. The system then compares metadata associated with the first version with metadata associated with the second version to create a difference report. Next, the system uses the different report to create an action plan that specifies how to merge the metadata associated with the first and second versions of the metadata object.

Correspondence Address:

**A. RICHARD PARK, REG. NO. 41241
PARK, VAUGHAN & FLEMING LLP
2820 FIFTH STREET
DAVIS, CA 95616 (US)**

(21) Appl. No.: **10/749,037**

(22) Filed: **Dec. 29, 2003**






FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5

DIFFERENCE REPORT 208

	TARGET	SOURCE	ACTION PLAN INPUT
TABLE	NAME	NAME	UPDATE
COLUMNS	COL 1	COL 2	UPDATE
	COL 2	COL 1	UPDATE
PROPERTIES		COL 3	CREATE
	PROPERTY 1		DELETE
		PROPERTY 2	CREATE
		PROPERTY 3	CREATE

FIG. 6

ACTION PLAN 404

702	704	706
<input type="checkbox"/>	TABLE 502 FROM TABLE 508	UPDATE
<input type="checkbox"/>	COLUMN 1 FROM COLUMN 2	UPDATE
<input type="checkbox"/>	COLUMN 2 FROM COLUMN 1	UPDATE
<input type="checkbox"/>	COLUMN 3	CREATE
<input type="checkbox"/>	PROPERTY 1	DELETE
<input type="checkbox"/>	PROPERTY 2	CREATE
<input type="checkbox"/>	PROPERTY 3	CREATE

FIG. 7




FIG. 8

**CUSTOMIZABLE METADATA MERGING
FRAMEWORK**

BACKGROUND

[0001] 1. Field of the Invention

[0002] The present invention relates to metadata that is used to define objects in databases. More specifically, the present invention relates to a method and an apparatus that facilitates merging metadata associated with different versions of objects of a database in a customizable manner.

[0003] 2. Related Art

[0004] Databases typically form the backbone of modern electronic business solutions. As business requirements change over time or are modified to meet local requirements, corresponding modifications may need to occur in the structure of an underlying database. These modifications can cause different versions of the same database to diverge over time. This can cause problems when divergent versions of a database need to be subsequently reconciled or updated.

[0005] Database objects are typically defined using metadata. The properties, attributes, and relationship of a database object with other database objects can be described using the unified modeling language (UML). The structure of metadata can be modeled using UML (Unified Modeling Language) to define a series of classes, associations, and attributes. For instance, a Table is a Class, a Column is a Class, and a Partition is a Class. 'Name' is an Attribute that exists for a Table, Column, and a Partition Class. A relationship exists between Table and Column (a table owns a column), and so we say there is an Association between the Table Class and Column Class. An end user typically receives metadata that defines a database object, for example a table, from a solutions provider. After receiving this metadata, the end user can typically customize the metadata to change the structure of the metadata object. For example, an end user may customize the metadata to add an additional column to a table or to change a property of the table.

[0006] It is possible for different end users to customize the same metadata in different ways. For example, end users in a multi-national corporation located in different jurisdictions may customize metadata to allow a corresponding database to accommodate specific data items required by their local jurisdictions.

[0007] In some circumstances, it is desirable to merge different versions of metadata. For example, a solutions provider may generate a new version of the metadata to enhance the database object, or management may want the metadata for two or more databases to be merged. This merging process is presently performed manually, and consequently involves a significant amount of time and effort.

[0008] Hence, what is needed is a method and an apparatus that facilitates merging different versions of a database without the problems described above.

SUMMARY

[0009] One embodiment of the present invention provides a system that facilitates merging different versions of a database object. During operation, the system receives metadata associated with a first version of a database object and a second version of the database object. The system then

compares metadata associated with the first version with metadata associated with the second version to create a difference report. Next, the system uses the different report to create an action plan that specifies how to merge the metadata associated with the first and second versions of the metadata object.

[0010] In a variation of this embodiment, metadata associated with the database objects are represented in Unified Modeling Language (UML).

[0011] In a further variation, comparing the metadata associated with the first and second versions of the database object involves allowing a user to specify which associations of the database object needs to be compared.

[0012] In a further variation, comparing the metadata associated with the first and second versions of the database object involves allowing the user to specify how to compare the metadata for each association in the database object.

[0013] In a further variation, creating the action plan involves allowing a user to select which actions to take in merging the metadata in order to produce merged metadata with the desired properties and attributes.

[0014] In a further variation, creating the action plan involves examining the difference report to determine what actions to take in bringing the metadata associated with the first version and second versions of the database object into agreement.

[0015] In a variation of this embodiment, the metadata associated with the first and second versions of the database object can define Oracle database objects like Dimensions, Cubes, Tables, Views, Materialized views, Transformations, and their associated objects. An example of associated objects is that Column, Foreign Keys, and Unique Keys can be associated with table metadata definition. Oracle is a trademark or registered trademark of Oracle International Corporation.

[0016] In a variation of this embodiment, the action plan can specify a number of actions including creating, updating, and deleting of database objects.

BRIEF DESCRIPTION OF THE FIGURES

[0017] FIG. 1 illustrates a metadata exchange framework in accordance with an embodiment of the present invention.

[0018] FIG. 2 illustrates the process of creating a difference report from compatible objects in accordance with an embodiment of the present invention.

[0019] FIG. 3 illustrates the process of creating a difference report from non-compatible objects in accordance with an embodiment of the present invention.

[0020] FIG. 4 illustrates the process of creating an action plan and a merged object in accordance with an embodiment of the present invention.

[0021] FIG. 5 illustrates source and target objects in accordance with an embodiment of the present invention.

[0022] FIG. 6 illustrates a difference report in accordance with an embodiment of the present invention.

[0023] FIG. 7 illustrates an action plan in accordance with an embodiment of the present invention.

[0024] FIG. 8 presents a flowchart illustrating the process of merging metadata associated with a first object and a second object in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION

[0025] The following description is presented to enable any person skilled in the art to make and use the invention, and is provided in the context of a particular application and its requirements. Various modifications to the disclosed embodiments will be readily apparent to those skilled in the art, and the general principles defined herein may be applied to other embodiments and applications without departing from the spirit and scope of the present invention. Thus, the present invention is not intended to be limited to the embodiments shown, but is to be accorded the widest scope consistent with the principles and features disclosed herein.

[0026] The data structures and code described in this detailed description are typically stored on a computer readable storage medium, which may be any device or medium that can store code and/or data for use by a computer system. This includes, but is not limited to, magnetic and optical storage devices such as disk drives, magnetic tape, CDs (compact discs) and DVDs (digital versatile discs or digital video discs), and computer instruction signals embodied in a transmission medium (with or without a carrier wave upon which the signals are modulated). For example, the transmission medium may include a communications network, such as the Internet.

[0027] Metadata Exchange Strategy Framework

[0028] FIG. 1 illustrates a metadata exchange strategy framework 102 in accordance with an embodiment of the present invention. Metadata exchange strategy framework 102 includes metadata change manager 110, metadata exchange strategy 112, and execution engine 114. Metadata change manager 110 provides functions for creating difference reports, such as difference report 108, as described below in conjunction with FIGS. 2, 3, 5, and 8. Metadata exchange strategy 112 provides functions for creating an action plan, such as action plan 113, as described below in conjunction with FIG. 4, 6, 7, and 8. Execution engine 114 provides functions for executing the action plan as described below in conjunctions with FIGS. 4 and 8.

[0029] Metadata exchange strategy framework 102 accepts source object 104 and target object 106 and creates merged object 108. Source object 104, target object 106, and merged object 108 contain metadata describing respective database objects. Source object 104 and target object 108 are descriptions of the same database object that differ, possibly because their metadata definition may have been updated or customized by different users.

[0030] Source object 102 and target object 106 are typically Oracle database objects and not just tables and columns (i.e. cubes, transformations, and the like). Also source object 102 and target object 106 may include other objects such as files in a file system or records in a file that have been captured using UML and stored in the database during design time. This is typically achieved by having a table called "file" and a table called "record." The metadata of external table objects refers to the record metadata in the record table. These are not database objects, but still dictate

the structure of some database objects. For example the structure of an external table is dictated by the file and record that it is referring to. Hence, there exists a requirement that external tables have to be merged from the record during design time, although recorded is not a database object. Hence, this is not just restricted to database objects but also to others that are marked in UML and stored in the database as classes, attributes, and associations.

[0031] Difference Report

[0032] FIG. 2 illustrates the process of creating a difference report from compatible objects in accordance with an embodiment of the present invention. Metadata change manager 110 receives external tables 202 and 204 and creates a difference report 208, which describes the differences between external tables 202 and 204. For example, external tables 202 and 204 contain metadata describing the same object that differ possibly because one of the external tables is an updated version of the other external table, or because the external tables have been customized by different individuals.

[0033] Difference Report from Non-Compatible Objects

[0034] FIG. 3 illustrates the process of creating a difference report from non-compatible objects in accordance with an embodiment of the present invention. In the example illustrated in FIG. 3, a difference report between external table 202 and record 302 is desired. Note that external table 202 may have been originally created from record 302.

[0035] Adapter 304 first converts record 302 into compatible external table 306. Metadata change manager 110 then receives external tables 202 and 306 and creates a difference report 308, which describes the differences between external tables 202 and record 302, which are ultimately the differences between external table 202 and record 302.

[0036] Action Plan and Merged Object

[0037] FIG. 4 illustrates the process of creating an action plan and a merged object in accordance with an embodiment of the present invention. Metadata exchange strategy 112 receives a difference report 208 created by metadata change manager 110. During this process, metadata exchange strategy 112 is used to examine difference report 208 and to create an action plan 404 for merging the two objects used in creating difference report 208.

[0038] Next, execution engine 114 receives action plan 404. After a user selects which items of action 404 are to be performed, execution engine 114 creates merged object 108.

[0039] Creation of an action plan requires the user to specify if a merge or a replace plan is preferred. The replace plan means deleting everything from the target the does not have a match in the source, while the merge plan would not allow deletion of anything from the target.

[0040] After the action plan has been created, the user can customize the action plan to apply the specific actions that are desired. Applying the action plan causes creation, deletion, or updates of the target object. Also, the user has the ability to specify which attributes in the target that need to be updated in the event of create or update in the target. The user can also change the values to suit the required needs. This is determined at the time when the user wants to execute the created action plan.

[0041] Source and Target Objects

[0042] FIG. 5 illustrates source and target objects in accordance with an embodiment of the present invention. These source and target objects include metadata for table 508 and table 502, respectively, and originated from the same metadata. Source table 508 includes a description 509 of table 508. Additionally, source table 508 is associated with columns 512 and properties 510. Properties 510 include property 2 and property 3, and columns 512 include column 1, column 2, and column 3. Together, table 508, properties 510, and columns 512 describe a corresponding database table.

[0043] Target table 502 includes a description 503 of table 502. Additionally, target table 502 is associated with columns 506 and properties 504. Properties 504 includes property 1, and columns 506 include column 1 and column 2. Together, table 502, properties 504, and columns 506 describe a corresponding database table. Note that any mapping between columns 506 of table 502 and columns 512 of table 508 is possible. For example, column 1 of columns 506 and column 2 of columns 512 may refer to the same data, while column 2 of columns 506 and column 1 of columns 512 may also refer to the same data.

[0044] Difference Report

[0045] FIG. 6 illustrates a difference report 208 in accordance with an embodiment of the present invention. Note that there are many ways to compare different components of the metadata. In this example, tables are being compared by name, columns are being compared by position, and properties are being compared by name.

[0046] Difference report 208 includes data for the parameters being compared, in this example table name, column position, and properties. The comparison of the table names indicates that the target and source have the same name; therefore the input into the action plan is to update the table.

[0047] The comparison of the columns indicates that column 1 in the target includes the same data as column 2 in the source and that column 2 in the target includes the same data as column 1 in the source. The input to the action plan is to update the column descriptions. Since column 3 exists in the source but not in the destination, the input to the action plan is to create column 3 in the target.

[0048] The comparison of the properties indicates that property 1 exists in the target but not in the source, while properties 2 and 3 exist in the source but not in the destination. The action plan for the properties includes deleting property 1 and creating properties 2 and 3. These actions will bring the target into alignment with the source if the user selects all of the actions as described below in conjunction with FIG. 7.

[0049] Action Plan

[0050] FIG. 7 illustrates an action plan 404 in accordance with an embodiment of the present invention. Action plan 404 includes actions derived from difference report 208 described above in conjunction with FIG. 6. Each line of action plan 404 includes a checkbox 702, a description of the action 704, and an action 706. A user can select which actions are required by checking the checkbox in column 702. After the user checks the actions, execution engine 114 performs the checked actions.

[0051] Merging Metadata

[0052] FIG. 8 presents a flowchart illustrating the process of merging metadata associated with a first object and metadata associated with a second object in accordance with an embodiment of the present invention. The system starts when a first metadata associated with the first version of an object is received (step 802). Next, the system receives the second metadata associated with the second version of the object (step 804). Note that steps 802 and 804 can occur in the opposite order or simultaneously. If needed, the system converts the second metadata to be compatible with the first metadata (step 806).

[0053] After receiving the first metadata and the second metadata, the system examines the first metadata and the second metadata to create a difference report (step 808). The system then creates an action plan from the difference report (step 810).

[0054] Next, the system allows a user to select portions of the action plan to implement by checking checkboxes in the action plan (step 812). Finally, the system uses the selected actions to merge the first metadata and the second metadata to create merged metadata (step 814).

[0055] The foregoing descriptions of embodiments of the present invention have been presented for purposes of illustration and description only. They are not intended to be exhaustive or to limit the present invention to the forms disclosed. Accordingly, many modifications and variations will be apparent to practitioners skilled in the art. Additionally, the above disclosure is not intended to limit the present invention. The scope of the present invention is defined by the appended claims.

What is claimed is:

1. A method to facilitate merging different versions of a database object, comprising:

receiving metadata associated with a first version of the database object and a second version of the database object;

comparing metadata associated with the first version of the database object with metadata associated with the second version of the database object to create a difference report; and

creating an action plan from the difference report that specifies how to merge metadata associated with the first version of the database object with metadata associated with the second version of the database object.

2. The method of claim 1, wherein metadata associated with the first version and the second version of the database object are represented in Unified Modeling Language.

3. The method of claim 1, wherein comparing metadata associated with the first version and second version of the database object involves customizing which associations to compare.

4. The method of claim 1, wherein comparing metadata associated with the first version and second version of the database object involves customizing how to compare the first metadata and the second metadata.

5. The method of claim 1, wherein creating the action plan from the difference report involves allowing a user to select

which actions to take in merging metadata in order to produce merged metadata with desired properties and attributes.

6. The method of claim 5, wherein creating the action plan involves examining the difference report to determine what actions to take in bringing metadata associated with the first version and second versions of the database object into agreement.

7. The method of claim 5, wherein metadata associated with first and second versions of the metadata object can define database objects, wherein database objects include tables, columns, dimensions, cube, views, materialized views, and external tables.

8. The method of claim 1, wherein the action plan can specify a number of actions including creating, updating, and deleting database objects, and their properties.

9. A computer-readable storage medium storing instructions that when executed by a computer cause the computer to perform a method to facilitate merging different versions of a database object, the method comprising:

receiving metadata associated with a first version of the database object and a second version of the database object; comparing metadata associated with the first version of the database object with metadata associated with the second version of the database object to create a difference report; and

creating an action plan from the difference report that specifies how to merge metadata associated with the first version of the database object with metadata associated with the second version of the database object.

10. The computer-readable storage medium of claim 9, wherein metadata associated with the first version and the second version of the database object are represented in Unified Modeling Language.

11. The computer-readable storage medium of claim 9, wherein comparing metadata associated with the first version and second version of the database object involves customizing which associations to compare.

12. The computer-readable storage medium of claim 9, wherein comparing metadata associated with the first version and second version of the database object involves customizing how to compare the first metadata and the second metadata.

13. The computer-readable storage medium of claim 9, wherein executing the action plan from the difference report involves allowing a user to select which actions to take in merging metadata in order to produce merged metadata with desired properties and attributes.

14. The computer-readable storage medium of claim 13, wherein creating the action plan involves examining the difference report to determine what actions to take in bringing metadata associated with the first version and second versions of the database object into agreement.

15. The computer-readable storage medium of claim 13, wherein metadata associated with first and second versions of the metadata object can define database objects, wherein database objects include tables, columns, views, dimensions, and cubes.

16. The computer-readable storage medium of claim 9, wherein the action plan can specify a number of actions including creating, updating, and deleting database objects and their properties.

17. An apparatus to facilitate merging different versions of a database object, comprising:

a receiving mechanism configured to receive metadata associated with a first version of the database object and a second version of the database object;

a comparing mechanism configured to compare metadata associated with the first version of the database object with metadata associated with the second version of the database object to create a difference report; and

a creating mechanism configured to create an action plan from the difference report that specifies how to merge metadata associated with the first version of the database object with metadata associated with the second version of the database object.

18. The apparatus of claim 17, wherein metadata associated with the first version and the second version of the database object are represented in Unified Modeling Language.

19. The apparatus of claim 17, wherein comparing metadata associated with the first version and second version of the database object involves customizing which associations to compare.

20. The apparatus of claim 17, wherein comparing metadata associated with the first version and second version of the database object involves customizing how to compare the first metadata and the second metadata.

21. The apparatus of claim 17, wherein executing the action plan from the difference report involves allowing a user to select which actions to take in merging metadata in order to produce merged metadata with desired properties and attributes.

22. The apparatus of claim 21, wherein creating the action plan involves examining the difference report to determine what actions to take in bringing metadata associated with the first version and second versions of the database object into agreement.

23. The apparatus of claim 21, wherein metadata associated with first and second versions of the metadata object can define tables, columns, properties of tables, and properties of columns.

24. The apparatus of claim 17, wherein the action plan can specify a number of actions including creating, updating, and deleting tables, columns, or properties.

* * * * *